Electronic Version
Stylesheet Version v1.1.1

Claims

[c1] 1.A power saving cache comprising:
circuitry to dynamically reduce the logical size of the cache in order to save power.

[c2] 2.A power saving cache according to Claim 1, further comprising means for determining an optimal cache size for balancing power and performance.

[c3] 3.A power saving cache according to Claim 2, wherein the means for determining an optimal cache size includes hardware means.

[c4] 4.A power saving cache according to Claim 2, wherein the means for determining an optimal cache size includes software means.

[c5] 5.A power saving cache according to Claim 1, further comprising means for maintaining coherency of data in the cache as the size of the cache is altered.

[c6] 6.A power saving cache according to Claim 5, wherein some of the data in the cache is modified data, and the means for maintaining coherency includes means for handling said modified data.

[c7] 7.A power saving cache according to Claim 1, wherein the circuitry includes means for partitioning the cache in one of several ways to provide a desired configuration and granularity.

[c8] 8.A power saving cache according to Claim 1, wherein the circuitry includes means to power off sections of the cache.

[c9] 9.A power saving cache according to Claim 1, wherein the cache is a set associative cache with N-ways, and the circuitry includes means to partition the cache along said ways.

[c10] 10.A method of operating a power saving cache comprising:
using circuitry to dynamically reduce the logical size of the cache in order to save power.

[c11] 11.A method according to Claim 10, wherein the cache is a set associative cache including N-ways, and the step of using circuitry to dynamically reduce the logical size of the cache includes the step of using the circuitry to partition the cache along the ways.

[c12] 12.A method according to Claim 11, wherein each of said N ways is individually powered.

[c13] 13.A method according to Claim 10, wherein the cache includes data, and the method comprises the further step of maintaining integrity of the data as the size of the cache is altered.

[c14] 14.A method according to Claim 13, comprising the further step of powering off sections of the cache.

[c15] 15.A method according to Claim 14, wherein some of the data in the cache is modified data, and the step of maintaining integrity of the data includes the step of, before powering off one of the sections of the cache, saving any modified data in said one section of the cache.

[c16] 16.A method according to Claim 15, further comprising the step of determining an optimum size for the cache given a set of power and performance criteria, and wherein the step of using circuitry includes the step of using circuitry to reduce the size of the cache to said optimum size.

[c17] 17.A method according to Claim 16, wherein the step of determining an optimum size includes the step of using one of a predefined set of hardware techniques to determine said optimum size.

[c18] 18.A method according to Claim 16, wherein the cache is

used in a processor, and the method includes the further steps of:

running on the processor a cache size adjustment routine; and

feeding information to said routine to determine the optimum size of the cache.

[c19] 19. A method according to Claim 10, wherein the using step includes the step of partitioning the cache in one of a given number of ways to provide a desired configuration and granularity, said given number of ways comprising (i) equal sized partitions, and (ii) binary weighted with or without a constantly powered way.